

# Reputations and Games

Sampath Kannan

Department of Computer and Information Science

University of Pennsylvania

June 10 2010

ONR MURI Meeting

Report Documentation Page				Form Approved OMB No. 0704-0188	
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE <b>10 JUN 2010</b>		2. REPORT TYPE		3. DATES COVERED <b>00-00-2010 to 00-00-2010</b>	
4. TITLE AND SUBTITLE <b>Reputations and Games</b>				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) <b>University of Pennsylvania, Computer and Information Science, 3451 Walnut St, Philadelphia, PA, 19104</b>				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT <b>Approved for public release; distribution unlimited</b>					
13. SUPPLEMENTARY NOTES <b>MURI Review, June 2010. U.S. Government or Federal Rights License</b>					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT <b>Same as Report (SAR)</b>	18. NUMBER OF PAGES <b>9</b>	19a. NAME OF RESPONSIBLE PERSON
a. REPORT <b>unclassified</b>	b. ABSTRACT <b>unclassified</b>	c. THIS PAGE <b>unclassified</b>			

# Reputation – A game-theoretic view

- Need repeated (pair-wise) interactions between agents = repeated games
- What does an agent do in a game?  
She plays a **(mixed) strategy** ... which may change over time, depending on opponent, etc.
- Her “reputation” should be a function of this time-varying mixed strategy.

# Repeated vs One-Shot Game: Example

- Prisoner's Dilemma

	Defect	Cooperate
Defect	-6	-1
Cooperate	-9	-2

*Payoffs to row player;  
Symmetrically to column player*

- Game-theorists talk about “type” of player – defecting type or cooperating type
- Reputation: inferred type of a player based on repeated observation

# To study value of reputation manager...

- Consider a 2-player game:
- Assume player Alice plays Bob repeatedly
- **New twist: A does not know the payoff matrix**
- Instead she must balance **exploration** and **exploitation** to minimize **regret**
- **Exploration:** A seeks to learn new matrix entries
- **Exploitation:** A seeks to profit from known entries
- **Regret:** A's lost payoff compared to the situation where she knows the matrix to start with

# If A knew B's type ...

- In zero-sum games she can identify her optimal strategies and learn payoffs for them with very little regret -  $O(n)$  regret where  $n$  is the number of strategies available to her. (This is best possible!)
- If instead she doesn't know B's type, best we can do currently is  $O(n^2)$  regret. Probably can't be beaten in general.
- Thus in this model, a reputation manager makes a big difference!

# More ways RMs can help

- Use player's behavior to estimate distribution of their types and their utilities for various outcomes
- If we assume a Bayesian prior on types, then we can design mechanisms [Z. Huang and Bei] that
  - Cause players to tell the truth about their types
  - Produce approximately optimal social welfare in some important resource allocation problems
- Reputation managers can give us these Bayesian priors



# More on Zhiyi-Bei result

- Mechanism design in Bayesian model for many optimization problems:
  - Combinatorial resource allocation
  - Submodular maximization
- These problems are known to be difficult to solve even approximately in worst-case model
- But in the Bayesian model (realized by having a reputation manager) there are good approximate solutions





# Future Work

- Explore other possible definitions of reputation in the game-theoretic context
  - Altruism: Each player has an altruism parameter  $a$  that corresponds to their reputation
  - Player gets payoff which is their own +  $a * (\text{everybody else's})$
  - Player's reputation is  $a$  and must be discovered
- In games with multiple equilibria: reputation is a way of describing which equilibrium a player prefers.



# Conclusions

- Game-theoretic frameworks can be used in two stages of our Trust Management infrastructure:
  - Defining and computing reputations
  - Making decisions based on these reputations
- While problems are still challenging, intractable problems assuming a worst-case adversary, could become tractable in this game-theoretic setting
- To use this idea one needs a more complete set of models for adversaries seeking to optimize their own objective functions